ISCR

Institute for Scientific Computing Research

Workshop and Conference Reports

CASC Center for Applied Scientific Computing

May 2000

Synopsis of Workshop Events:

Power Programming Short Course

The Institute for Terascale Simulation, an ASCI-supported arm of the ISCR, organized a 3-day *Power Programming Short Course* to enable laboratory code developers to come to grips with performance-oriented aspects of parallel computation on distributed shared memory machines, such as the ASCI White machine that arrived at the laboratory in July 2000. The instructors were Steve White (IBM), Larry Carter (UCSD), David Culler (UCB), ISCR collaborator Clint Whalley (University of Tennessee), and Bill Gropp (Argonne National Laboratory). Sixty-five people (LLNL training center room capacity) attended—most for the full three days, May 15-17, 2000.

The topics were as follows:

- POWER3 Architecture and Tuning for the ASCI White system: the POWER3 microarchitecture, tuning advice and experience, and ASCI White overview (White)
- Performance Programming, I: exploiting the Power processor in the design of kernels and data structures for scientific applications (Carter)
- Performance Programming, II: cache and TLB issues (Carter) Sensitivity-based Performance Analysis Tools: understanding performance thresholds through parameterized microkernels (Culler)
- High Performance Communication, I: MPI-1, point-to-point and collective (Gropp)
- High Performance Communication, II: parallel I/O, MPI/OpenMP tradeoffs, and communication monitoring tools (Gropp)

The workshop was of value beyond the ASCI context, since the Power architecture is in many ways typical of other hybrid architecture high-end scientific computers. Participants were encouraged to bring laptops or overhead slides with which to display code or precisely state questions or comments.

The short course was organized by David Keyes of the ISCR, John May of CASC, and Mary Zosel of Livermore Computing.

For more information, see the website
http://www.llnl.gov/CASC/workshops/course_reg/short_course_info.html

Synopsis of Workshop Events:

The Second Workshop on Mining Scientific Datasets

The Second Workshop on Mining Scientific Datasets was held July 20-21, 2000, at the Army High Performance Computing Research Center (AHPCRC) in Minneapolis, Minnesota. The goal of the workshop was to bring together researchers from the data mining and scientific computing communities in order to better understand how data mining can be used for the exploration of scientific datasets. Whereas the literature of data mining is dominated by demographic, economic, and related applications, the participants concentrated on the challenges and opportunities that are unique to the application of data mining to scientific data sets.

This workshop followed an earlier workshop on the same topic that was held in September 1999 at AHPCRC. The first workshop covered a broad range of topics with the aim of identifying the main disciplines that could help address the challenges in mining scientific data sets. The second workshop focused on how the research being done in these diverse technical areas could be effectively harnessed to solve the problems of scientific data analysis.

The 22 talks at the workshop were broadly divided into two categories:

- Scientific applications where data mining techniques were either being applied, or could be applied. These included areas such as astronomy, physics, earth sciences, fluid dynamics, protein folding, and spatial data sets arising in Geographic Information Systems.
- Data mining algorithms, including traditional pattern recognition techniques such as clustering, classification, and association rules, as well as various data pre-processing techniques such as dimension reduction, feature extraction, and feature selection.

A few speakers gave a broad overview of the field and shared their experiences with mining scientific data. Jagdish Chandra, Professor at George Washington University and former director of the Math and Computer Science Division at the Army Research Office, described the mathematical challenges in understanding high-dimensional data sets. He identified several common features of current data sets, including their massive size (terabytes and petabytes), their complex structure in terms of the relations between different parts of the data, their noisiness as a result of the way in which the data was collected, as well as our imperfect understanding of the basic processes for which we seek information.

Chandra observed that questions regarding the structure of the data had been present for centuries. However, it was the scale of the problem, as well as its internal complexity, that required us to rethink old ideas and solve the problem with an infusion of new ideas. He characterized these “model rich” data mining problems by their non-linearity as well as intrinsic uncertainty. He believed that computational power by itself was not enough to solve these problems—a deep understanding of the natural and mathematical structure underlying the new data was necessary in order to extract information with scientific, technological, and societal implications. He echoed the views of several experienced data miners when he observed that, to address these problems, we need a multi-disciplinary approach, combining the expertise of mathematicians, numerical analysts, signal processing engineers, statisticians, and computer scientists.

Synopsis of Workshop Events (continued):

A different perspective was given by Padhraic Smyth from University of California at Irvine, who described some of the new research ideas being pursued in data mining. He observed that data mining could be considered as the data driven discovery of models and patterns from massive observational data sets. This would require a modern-day data miner to be proficient in statistics, modeling, optimization, and search, as well as algorithm design and data management. The process of data mining could be considered as being composed of several components, including the dataset itself, the task (e.g., prediction, clustering, etc.), the representation of the problem (i.e., the underlying language used), the score function and the optimization technique used, and the data access approach. Each of these components would be driven by the application. Smyth identified three areas of new research currently being pursued in data mining—scalability either by scaling down the data or scaling up the algorithms, the identification of patterns vs. models, and the clustering of objects of different dimensions.

A broad range of application areas was covered by the workshop participants. The problems being tackled by means of data mining techniques included the detection of coherent structures in turbulent boundary layers, the prediction of three-dimensional contact potentials among protein residues, the identification of galaxies with a bent-double morphology, the content-based querying of earth science data, as well as the Virtual Observatory concept being proposed by astronomers as they mine several large-area sky surveys. In the area of algorithms, the topics included the generation of long patterns, dimension reduction of image feature descriptors, clustering via the construction of a decision tree, as well as techniques and protocols for mining distributed data.

The workshop included a panel discussion on spatial data mining. The panelists, led by Sashi Shekhar (University of Minnesota), included Jiawei Han (Simon Fraser University), James LeSage (University of Toledo), Suchi Gopal (Boston University), and Sanjay Chawla (Vignette Corporation). The discussion centered around the challenges encountered in mining spatial data, including the appropriate ways to discretize a spatial domain, the selection of features in light of the diversity of possible spatial relationships, the metrics for evaluation that capture spatial accurately, etc.

During the workshop, several talks touched upon the problems often encountered in mining scientific datasets, including:

- the heterogeneity of the data
- the high dimensionality of the feature space
- the problems in extracting relevant features from simulation data
- the spatio-temporal nature of the data
- feature selection
- scaling the algorithms to massive data sets
- the presence of noise in the data
- the lack of ground truth

Synopsis of Workshop Events (continued):

Dr. N. Radhakrishnan, the Director of the Corporate Information and Computing Directorate of the Army Research Laboratory, closed the workshop with encouraging words for the participants, urging them to address the challenges that remain in applying data mining techniques to massive scientific data sets resulting from computer simulations and observations.

The overwhelming interest expressed by the 110 attendees of this second workshop has led to a follow-on workshop that will be held in conjunction with the first SIAM International Conference on Data Mining. Additional details on this conference are available at *http://www.siam.org/meetings/sdm01/*.

The workshop was sponsored by AHPCRC, the Minnesota Supercomputing Institute, and the Center for Applied Scientific Computing at the Lawrence Livermore National Laboratory. The co-organizers were Bob Grossman (University of Chicago and Magnify, Inc.), Chandrika Kamath (Lawrence Livermore National Laboratory), Vipin Kumar (University of Minnesota), and Raju Namburu (Army Research Laboratory).

A website with details on the first and second workshops, together with abstracts and presentation slides, is at
http://www.ahpcrc.org/conferences/

Contributed by Chandrika Kamath, Center for Applied Scientific Computing, LLNL.

Workshop on Solution Methods for Large-Scale Nonlinear Problems

Synopsis of Workshop Events:

The solution of large-scale, fully coupled multiphysics models is vital to progress in computational science. Efficiently computing solutions to the mathematical equations underlying these models requires effective numerical algorithms for solving very large systems of coupled nonlinear equations. Scalable algorithms that can exploit the power of massively parallel computers are especially necessary. Several successful approaches to constructing such algorithms have emerged in recent years, but the need to treat more complex models and problematic physics continues to drive research.

The Workshop on Solution Methods for Large-Scale Nonlinear Problems, held July 26–28, 2000, at the Four Points Hotel in Pleasanton, CA, brought together many of the most active researchers in nonlinear solution algorithms and applications. The workshop was hosted by the Center for Applied Scientific Computing and the Institute for Terascale Simulation at Lawrence Livermore National Laboratory (LLNL). Major themes included Newton–Krylov methods, nonlinear multigrid methods, preconditioning techniques, operator-split and fully implicit schemes, and large-scale sensitivity analysis and optimization. Major application areas included general fluid dynamics, combustion and other reacting flows, porous media flows, radiation diffusion problems, and design optimization.

The workshop was a distant follow-on to similar workshops on large-scale nonlinear problems held at Utah State University in 1989 and 1995. In the spirit of those earlier workshops, the schedule of talks and discussion sessions allowed liberal break time to encourage informal interactions among the participants. Workshop attendees came from academia (25), government laboratories (21), and industry (1) and were mainly from the United States.

Each of the three workshop days emphasized particular subject areas, with nine presented talks followed by a moderated discussion of selected issues. The first day focused on Newton–Krylov methods and nonlinear multigrid, also called the Full Approximation Scheme (FAS); the second day focused on applications and problem formulation; and the third day focused on algorithmic issues, sensitivity analysis, and optimization.

The first day began with two talks on techniques for enhancing the convergence of Newton–Krylov methods. The first introduced a novel nonlinear preconditioning method based on Schwarz domain decomposition; the second outlined approaches based on the singular value decomposition. Later presentations addressed applications of Newton–Krylov and nonlinear multigrid methods to combustion, groundwater flow, and radiation diffusion problems. The final talk addressed the benefits of using automatic differentiation instead of finite differences in computing matrix-vector products in Newton–Krylov methods. To conclude the day, Van Henson and Jim Jones led a moderated discussion on the differences in performance between Newton's method and FAS on various problems. The basic iteration scheme of FAS was discussed along with the full multigrid version. Methods using multigrid to solve Jacobian systems for each Newton iteration (Newton-multigrid) were presented. The consensus of attendees who had experience with FAS indicated that FAS requires less memory and has a larger basin of attraction than Newton-multigrid methods, but that it is harder to get FAS to work on new problems.

Synopsis of Workshop Events (continued):

The second day, which emphasized applications and problem formulation, included presentations on nonlinear solvers for reacting flows, multiphase groundwater flow, chemical reactions, radiation diffusion problems, powder consolidation, and Einstein's equations of general relativity. Solution approaches included operator splitting techniques and fully implicit formulations, domain decomposition, and pseudo-transient continuation. Peter Brown and Carol Woodward led the moderated discussion that afternoon on whether or not to operator-split coupled systems of nonlinear equations. Participants generally agreed that the holy grail is a fully implicit solve of all problem physics, but simulation technology has not evolved enough to handle all the necessary complexity. Significant progress has been made, however, in development of solvers and methods for these fully implicit formulations.

The third day, which focused on algorithmic issues, sensitivity analysis, and optimization, included presentations on nonlinear elimination methods, performance "stress points" of parallel implicit solvers, techniques for finding solution sensitivities to problem parameters, and formulation and solution techniques for nonlinear PDE-constrained optimization. Applications included aerodynamic analysis and design optimization as well as optimal control of Navier–Stokes flows. Steven Lee and David Keyes moderated the afternoon discussion on the role of sensitivity analysis and optimization in scientific computing. Discussion participants pointed out that aerodynamic and automotive design were fields where sensitivity analysis and optimization were now being used routinely and that, with the development of effective automatic differentiation tools, more fields will begin using these techniques in the future.

From the discussion and presentations, it became clear that the field of solution methods for large-scale nonlinear problems has evolved considerably in the five years since the last workshop. Although Newton–Krylov and Newton–Krylov–Schwarz methods were important topics at the 1995 meeting, they were much more prominent in this workshop. This change reflects advancement of the Newton–Krylov family of methods and the current widespread acceptance of these techniques as the methods of choice in many difficult applications. Similarly, fully implicit and fully coupled problem formulations were much more in evidence in this workshop as a result of advances in solution technology that have made fully coupled formulations feasible in previously intractable applications. Such advances have also made it possible to more effectively treat difficult PDE-related optimization problems, such as design optimization. Accordingly, the optimization talks in this workshop focused more on PDE-related applications and less on general optimization algorithms.

A notable new feature of this workshop was the work on nonlinear multigrid, which did not appear at all in the previous workshop. Although the

Synopsis of Workshop Events (continued):

method is not new, it has not been widely used until recently, when advances in the solution of heterogeneous problems made the method more applicable to problems of interest. Another area not previously featured was sensitivity analysis. Even though this area has been a subject of importance for some time, the introduction of special Newton–Krylov techniques in recent years has resulted in much more effective methods for large-scale problems.

Many of the workshop topics seem likely to lead to interesting future research. Nonlinear preconditioning, which generated much discussion among the participants, shows considerable potential for resolving well-known stagnation problems associated with Newton's method. Nonlinear multigrid, though not as new, also appears to have much unexploited potential, not only as a nonlinear solver but perhaps as a nonlinear preconditioner as well. Similarly, sensitivity analysis and design optimization seem likely candidates for further attention as algorithms are improved and more challenging applications addressed. Another area, the use of operator splitting techniques for preconditioning fully implicit formulations, should lead to more accurate simulations using fully coupled approaches. This technique will also allow the reuse of current methods and codes to solve these preconditioning systems. Nonlinear elimination, which was observed to be useful in treating shocks in compressible flow applications, also appears worthy of further attention. Finally, automatic differentiation seems at last to have emerged as a practical tool that should see many applications in Newton–Krylov methods. Automatic differentiation has the potential to alleviate algorithmic slowdowns and failures that sometimes occur when finite-difference techniques are used in approximating matrix-vector products.

A special issue of the journal, *Numerical Linear Algebra with Applications*, will feature papers from this workshop. The issue will be co-edited by workshop organizing committee chair Carol Woodward and journal editor Panayot Vassilevski. On the basis of the enthusiasm of workshop participants, the apparent algorithmic advances, and the challenge of increasing problem scale, LLNL's Center for Applied Scientific Computing and Institute for Terascale Simulation will sponsor more workshops in this series in coming years.

For more information, see the website
http://www.llnl.gov/CASC/workshops/workshop_reg/workshop072600_info.html

Contributed by Homer Walker, Mathematical Sciences Department, Worcester Polytechnic Institute and Carol S. Woodward, Center for Applied Scientific Computing, LLNL.

Synopsis of Workshop Events:

Computational Science Graduate Fellows Conference

The Krell Institute, which operates the *Computational Science Graduate Fellowship* (CSGF) program for the Department of Energy, organizes a summer research symposium for fellows and their advisors from leading universities across the United States. In 2000, for the first time, Krell decided to host the symposium on site at the three DOE research laboratories in the San Francisco Bay area. The ISCR assisted Krell with the local logistics of the three-day conference, which included on-site presentations at the Lawrence Livermore, Sandia Livermore, and Lawrence Berkeley laboratories. Speakers were provided by all three of the local laboratories, and other DOE organizations.

Seven CSGF students were interning at the ISCR already at the time of the July 27–29 meeting. They were joined by thirty-nine of their peers and a contingent of ten advisors. Thirty-six DOE scientists from all over the national Office of Science and Defense Programs laboratory complex, and DOE administrative staff from Washington, D.C., converged at the conference to provide technical content and practical career advice to the “cream of the crop” graduate students.

David L. Brown of CASC gave one of the plenary presentations at the Symposium, entitled *Overture: An Object-oriented Framework for Solving Partial Differential Equations in Complex Geometry.*

Evi Dube of B Division described *Computational Research Challenges for an ASCI Simulation Code* in one of the breakout sessions.

Another major LLNL code project, *Using AMR in CFD – from Shock Tubes to Lasers*, was presented by A Division’s Jeff Greenough.

Richard Hornung of CASC presented a faculty collaborative ISCR project entitled *A Hybrid Model for Gas Dynamics that Couples Continuum and Direct Simulation Monte Carlo Methods Using Adaptive Mesh Refinement* at another LLNL breakout.

Completing the LLNL presentations was a student collaborative ISCR project, *The ROSE Project: The Optimization of Object-oriented Scientific Applications*, presented by Dan Quinlan of CASC.

For more information about the CSGF program, see the website at *http://www.krellinst.org/CSGF/*

The conference website is *http://www.krellinst.org/CSGF/conference.html*

Synopsis of Workshop Events:

Internships in Terascale Simulation Technology

In collaboration with LLNL's Science and Technology Education Program (STEP), and with sponsorship from the DOE Defense Programs Office in Washington through Beverly Berger, the ISCR organized and conducted a series of tutorial lectures to enrich the technical and social experiences of the large population of summer students interning at the laboratory. Dubbed *Internships in Terascale Simulation Technology* (ITST), this ten-week series with two lectures per week featured ten different presenters.

Led and anchored by Alice Koniges, fresh from the success of her new Morgan-Kaufman book *Industrial-strength Parallel Computing*, with six lectures, the tutorial also featured two sessions each from two other recent LLNL authors: John May, whose tome on *Parallel I/O* (also from Morgan-Kaufman) appeared in September, and Van Emden Henson, who co-authored an update of the 1987 SIAM classic, *A Multigrid Tutorial*.

Five LLNL computational physicists shared from their experiences with real-world applications related to the ASCI mission. David Brown presented LLNL's *Overture* computational environment for the solution of PDEs in complex geometry, featuring several examples from combustion, immiscible fluids, and other fields. Garry Rodrigue gave a two-lecture introduction to the numerical analysis of shock physics. Howard Scott, Alek Shestakov, and Lin Yang combined for two lectures on parallel programming techniques and data communication paradigms in practical physics codes. David Keyes, who together with Koniges organized the overall program, finished up with two lectures on high-performance parallel algorithms for PDE simulation, which included an "anatomy" of the Bell-prize-winning computations he performed with his graduate student Dinesh Kaushik and collaborators from DOE's Argonne National Laboratory and NASA.

Recently a student himself, CASC's Gary Kumfert gave a workshop on the practical matter of preparing scientific presentations. This topic was presented at the summer's midpointæafter the interns were sufficiently into their projects to think concretely in their terms but with plenty of time left before STEP's Poster Presentation day in August, at which several ISCR students showed off the fruit of their labors.

It was gratifying to find permanent CASC researchers attending an occasional subseries. For example, several CASC computer scientists attended Garry Rodrigue's lectures on numerical simulation of shocks after finding, being employed at the laboratory for a season, that they needed to understand the associated vocabulary and challenges to converse with laboratory "clients." Similarly, veteran computational researchers took advantage of John May's tutorial lectures on parallel debugging and parallel I/O.

Synopsis of Workshop Events (continued):

Schedule of Lectures

June 13

Parallel Computing Resources/Parallel Architecture Overview
Speaker: Alice Koniges

June 15

Performance Issues, Measuring and Reporting Performance
Speaker: Alice Koniges

June 20

Parallel Programming Models and Languages I
Speaker: Alice Koniges

June 22

Programming Models and Languages II, Parallel I/O, Parallel file systems
Speaker: Alice Koniges

June 27

Performance Optimization, Optimization Issues
Speaker: Alice Koniges

June 29

Case Studies: How Much Can Performance be increased in a Real Application?
Speaker: Alice Koniges

July 6

Power Presentations
Speaker: Gary Kumfert

July 11

Basic Parallelization,
Speaker: Alek Shestakov
and
Parallel Applications in Physics,
Speaker: Lin Yang

July 13

Mixed Models (Open MP or Pthreads with MPI),
Speaker: Howard Scott

Synopsis of Workshop Events (continued):

July 18

PDEs in Complex Geometry
Speaker: David Brown

July 20

PDEs in Complex Geometry
Speaker: David Brown

July 25

Numerical Shock Simulation
Speaker: Garry Rodrigue

July 27

Numerical Shock Simulation
Speaker: Garry Rodrigue

August 1

Parallel I/O and Parallel Debugging
Speaker: John May

August 3

Parallel I/O and Parallel Debugging
Speaker: John May

August 8

A Multigrid Tutorial
Speaker: Van Henson

August 10

A Multigrid Tutorial
Speaker: Van Henson

August 15

Parallel Solver Infrastructure
Speaker: David Keyes

August 17

Parallel Solver Infrastructure
Speaker: David Keyes

Synopsis of Workshop Events:

Fifth Symposium on Overset Grids & Solution Technology

Overset grid technology has emerged as an effective method for providing discretizations of complex geometries for physical simulations. Instead of creating a single grid to represent a geometry and the domain of interest, a collection of overlapping, or overset, component grids is provided; each grid represents a portion of the domain. Solving a system of equations over the entire problem involves determining a solution on each component grid and communicating the solution between grids in overlapping regions.

As the latest in a ten-year series of semiannual symposia, the Fifth Symposium on Overset Grids & Solution Technology provided an effective forum for ideas and applications. The symposium was hosted by the University of California at Davis and held at that campus on September 18th–20th, 2000. A diverse range of institutions further supported the meeting representing government, academic institutions, and industry. Sponsors included the Institute for Scientific Computing Research (ISCR) at LLNL, the U.S. Army Research Office, UC Davis Department of Mechanical & Aeronautical Engineering, and Intelligent Light Corporation, the developers of a popular engineering visualization environment.

Presentations and discussions during the symposium involved a variety of overset grid topics spanning numerical methods, grid generation, applications, and software development. The meeting was organized to provide a mixture of these topics each day, concluding with a panel discussion entitled “Pros and Cons of Overset Grid Solution Technology.” Participants represented government, academic, and industrial researchers from around the world, applying overset grids to many different applications. A diverse set of application domains were presented including interface dynamics, moving body simulations, chemical vapor deposition, tidal flow simulation, and unsteady insect flight dynamics.

Discussions regarding numerical methods were as diverse as the applications they were intended to model. Brian Miller and Bill Henshaw, members from the Overture project at Lawrence Livermore’s Center for Applied Scientific Computing (CASC), presented their development of level set methods on overlapping grids. Dan Quinlan and Bobby Philip, also from the Overture project, discussed their application of overset grids in a hierarchical manner to build adaptive mesh elliptic equation solvers. ISCR summer student Lars Carlson from Chalmers University in Sweden related his recent work on line-implicit methods for the incompressible Navier-Stokes equations, a topic made more difficult by the unstructured nature of overlapping grid connectivity. Robert Trammel, of CFD Research Corporation, presented papers on time integration techniques for unsteady flow applications and overlapping additive Schwarz methods for the Helmholtz equations. On the second day, Anders Petersson, an Overture team member, provided an investigation of pressure boundary conditions for the incompressible Navier–Stokes equations.

Grid generation remains an area of considerable activity within the overset grid community. Many of the mesh generation presentations focused on running larger, more complicated problems and devising methods to automate the generation of the collection of meshes. William Chan from NASA Ames presented two papers highlighting recent developments of his interactive mesh generator OVERGRID and his automated surface mesh generation technologies. There were two talks describing algorithms for the automated assembly of the overlapping meshes, including the creation of the necessary interpolation stencils. Representatives from industry demonstrated the capabilities of their products.

Synopsis of Workshop Events (continued):

Meng-Sing Liou, a researcher at NASA Glenn, described an interesting alternative to overlapping component meshes using a new technology that replaces the overlapping regions in an overset grid with unstructured patches. An unusual and impressive demonstration of the overlapping mesh concept was provided by Professor Nakahashi of Tohoku University, Japan, who used overset unstructured meshes for large moving body problems.

Many applications of overset grid technologies were represented at the symposium. One group utilized overlapping grids to investigate three-dimensional unsteady flow of blood in major arteries. Another researcher used the Overture software framework to simulate chemical vapor deposition. The first day saw Petri Fast of the Overture project present a method for modeling interface dynamics using overset grids. Aerodynamic applications included wind power generation, modeling the interactions of bodies in buoyant jets and plumes, and turbomachinery simulations for the space shuttle. One notable example was provided by Dora Yen, a UC Davis graduate student, modeling the use of micro-electro-mechanical (MEMs) devices to control flow over wing sections. Her work coupled the use of wind tunnel experimentation and computer simulation using overset grids to investigate this new method for controlling the flow around aircraft wings. Overset grids are often used to simulate bodies moving relative to one another within a flow field. Several examples of this class of problem were presented at the symposium, including simulations of store separation from aircraft, space shuttle booster launch modeling and the deployment of multi-warhead projectiles.

Several presentations described improved software techniques making the overset grid method more efficient and accessible to other researchers. A topic permeating many talks was the exploitation of parallel computer architectures within existing overlapping grid tools. Brian Gunney presented Overture's recent developments in this area while Cetin Kiris described NASA Ames' recent efforts. Software framework projects were also important contributors to the symposium. The Overture Project from CASC produced several talks, both from within the group itself as well as from users. Andrew Wissink, another CASC representative, presented related technology in the SAMRAI adaptive mesh framework. Variations and applications of NASA's OVERFLOW code were also prominent. Mississippi State University researchers presented another software toolkit for overlapping mesh problems. As a proponent of the technology, the Army Research Labs proposed to include overset grid software as a part of their Major Shared Resource Center.

The symposium concluded with panel discussion on "The Pros and Cons of Overset Grid Solution Technology." As evidenced by the many and varied applications presented, most panelists were already in agreement that the method was flexible and effective. Current research issues such as automated mesh generation and parallel computing were emphasized as important for the continued development of the technology. Several expressed the desire to see the techniques used in an even broader community of researchers and engineers, encouraging the development of freely available software to accomplish this end. Concluding the remarks as well as the symposium, participants received an invitation to the Sixth Overset Grid Technology Symposium, to be held in the Washington, DC, area in 2002.

For more information, see the website *http://ntserver.itd.ucdavis.edu/Chimera2000/*

Contributed by Kyle Chand, CASC

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551